# National Aeronautics and Space Administration (NASA)
# Earth Observing System Data Information System (EOSDIS)

## Case Study

John Hrastar

January 2003

Table of Contents



This study is sponsored by the NASA Academy of Program and Project Leadership (APPL), Code FT NASA Headquarters, and the Systems management Office at the Goddard Space Flight Center.

## EXECUTIVE SUMMARY

The Earth Observing System Data and Information System (EOSDIS), one of NASA's largest, most complex projects, was started as part of the Earth Observing System (EOS). This system was meant to collect, process, distribute, and archive the large amount of data that was to be generated by the EOS program and to archive and distribute NASA Earth science data .

The early concepts for the EOSDIS were shaped by NASA experiences with previous data systems. A large centralized data system was necessary to insure on-time data processing, as well as distribution to those involved in the interdisciplinary research of Earth science. Data could not be hoarded by the Principal Investigators as had been previously done. A large aerospace contractor was seen as the ideal developer, because these contractors had experience delivering large systems to the Department of Defense (DoD). In general, there was little experience within NASA or other civilian agencies with data systems of this magnitude and complexity.

During the start-up phase, contractor phase B studies were done and requirements were generated. However, these were done in the context of traditional procurement. A contract was signed for the development of a large, centralized EOSDIS Core System (ECS) to meet a fixed-end item specification. The issue of a continuing, long-term (10 year) development plan that evolved with technology was recognized, but this issue was never fully resolved with the work in the contract. Some involved saw it necessary to put blinders on and build the EOSDIS to the specifications. Others saw it as an R&D system that must—and could—evolve with time.

The split between the developers and the users came early in the implementation phase, and it was driven by these different views of the system. In the early to mid 90's, attempts were made to reconcile these views and change the system to accommodate everyone. Such modifications, coupled with problems of poor performance, caused delays and budget overruns.

Simultaneously, external pressures were also having an effect. EOSDIS budgets were significantly reduced to save the Agency and the Office of Earth Science money for other purposes. This external reduction in funding was dramatic, reducing EOS funding by approximately 10 Billion dollars by mid-decade. The impact on the EOSDIS was significant. Most of the time, these reductions were not accompanied by corresponding reductions in requirements. Project visibility and poor performance combined to encourage reviews by many prestigious external review teams.

Under changing leadership in the mid to late 90's, the necessity for basic changes was recognized to avoid total collapse of the EOSDIS. The focus was changed from a centralized system being developed all at once, to a smaller more doable system that would be developed incrementally. The main contract requirements were cut back to a core to be developed by the contractor, the science teams developed innovative ways of doing the science data processing. Thus, the system became more distributed.

Under these revisions, the EOSDIS was completed and is now delivering data to users. It does not have all the capabilities originally requested, but it also did not cost as much as originally budgeted. It handles over a terabyte of data per day, and includes the management of several petabytes of data, over 1.75 million lines of code, and support of over 900 data sets. Nevertheless, it is an impressive achievement that serves more than 2 million users per year.

A number of lessons can be drawn from the difficult path followed by the EOSDIS during development.

## INTRODUCTION

The Earth Observing System Data and Information System (EOSDIS) is one of NASA's largest, most complex systems. It was conceived and built as the control, collection, and processing system for all Earth science data collected as part of the Earth Observing System (EOS) program and as the distribution and archiving system for the Earth Science Enterprise of NASA. It was also to include the mission command and control of the EOS flight missions. The EOSDIS differs from the preceding NASA data systems both in size and the interoperability concept. NASA had extensive experience in developing science information systems for individual missions, but a multi-mission system with data volumes as large as those produced by EOS had never been attempted. It handles more than a terabyte of data per day and includes the management of several petabytes of data, over 1.75 million lines of code, and support of over 900 datasets. The EOSDIS serves more than 2 million users each year, many of them non-scientists. This is significantly greater than the 10,000 science users originally envisioned. It is also a system significantly different from other complex NASA systems like Apollo, the Shuttle, and the Space Station. Unlike these systems, EOSDIS is not a flight system, and it is therefore outside the traditional expertise of the Agency.

> The major challenge for the EOSDIS was to achieve a balance between the mission-critical operational activities and the necessary progress in science and technology.

The former requires a stable, robust system; the latter continued evolution, innovation, and change. Striking such a balance is the unique feature of an R&D environment as compared with a pure operational environment found in other Agencies. This challenge is not unique to EOSDIS, but extends to future systems as well.

The EOS program was conceived in the 1980's as an interdisciplinary approach to studying the Earth. This would require many remote sensing instruments on a number of large spacecraft to repeatedly scan the environment and develop a 15 year (or more) data set. The need for the 15 year data set was driven by the desire to include one complete solar cycle. This data set would be available to the Earth science community to help characterize Earth and its changes. The results would not only increase this type of understanding, but it would provide information to those outside the science community (e.g. policymakers) to assure sustainable development on Earth.

The EOSDIS concept was developed with EOS. Its program structure was partly a reaction to some of the problems experienced on the missions of its predecessors, and partly the recognition that an interdisciplinary approach involving multiple science communities (atmospheres, oceans, land) would require a different kind of system than had previously existed. The EOSIDS is also NASA's contribution to the Global Change Data Information System (GCDIS).

The EOSDIS was developed in the 90's with the EOS flight program, and it is now operating with data from many spacecraft. The development, however, was anything but smooth. In addition to the environmental changes experienced by the Agency in the 1990's, the EOSDIS experienced its own development difficulties, some of which were self-imposed.

The purpose of this case study is to help NASA managers, engineers, and scientists understand what happened during the implementation of the EOSDIS in order to apply the lessons learned to future programs and projects. This study focuses on lessons learned--both successes and failures—and is not intended to give a detailed history. There are neither detailed traces of the Level 1 requirements, nor traces of the Program Operating Plan (POP) histories included in this report.

Information in this case study was gathered mainly from principals involved with the development of EOSDIS, and from reference material. The EOSDIS development history is a long one, beginning in the early 80's and ending in 1999. This history includes many internal and external changes. Though generally consistent, some accounts taken from the memories of those involved are at odds with one another. While it is unlikely that these events will be repeated, a review of the EOSDIS development provides a number of lessons learned for future program managers.

## EOSDIS Structure

The EOSDIS is part of the EOS ground system. Its structure has evolved over the years from the original concept, including the renaming of some of the components. The following is a description of the present system.

The EOSDIS is divided into six major components[1]: data capture by the Tracking and Data Relay Satellite system (TDRSS) and EOS Polar Ground System (EPGS); the EOS Data and Operations System (EDOS); the EOSDIS Backbone Network (EBNET); the EOSDIS Core System (ECS); the Distributed Active Archive Centers (DAAC); and the Science Investigator-led Processing Systems (SIPS). The TDRSS and the EPGS capture spacecraft science data, forward it to the EDOS for processing, and provide the telemetry and command link for controlling spacecraft health and safety. The EDOS separates the data by instrument and performs the initial processing and back-up archiving. The EBNET delivers real-time data to and from the operations control centers, and delivers the science data to the DAACs. The ECS provides satellite and instrument command and

[1] For detailed overview of these components see: http://esdis-it.gsfc.nasa.gov/add/overview.htm

control, as well as data product generation, archival, and distribution. Product generation is done using science software provided by the Principal Investigators (PI).

Using the ECS, the DAACs process the raw data into useful products, handle all user product searches, requests, and orders, and distribute data and information directly to the user community via the internet. The DAACs also archive the data and information. Each DAAC focuses on the data needs of a specific segment of the user community. Instrument teams may propose to produce their standard products operationally and deliver them to the DAACs instead of having the DAACs do this. This would be done in the SIPS. The decision to use SIPS is made on a case-by-case basis. The introduction of the SIPS was a late development in response to ECS development problems and will be discussed later in the report. Because most of the lessons learned dealt primarily with the ECS, this report will concentrate on that area.

## EARLY CONCEPTS

Experiences on previous and contemporary projects helped to shape the ideas for EOS and the EOSDIS. In the early 80's, a NASA Associate Administrator (AA) realized that the crowded geosynchronous orbit would eventually require a larger communications spacecraft to accommodate the increasing communications traffic. Although not meant for geosynchronous orbits, this idea was extended by the AA to apply to Earth science spacecraft (i.e. large multi-instrument platforms instead of many smaller spacecraft.)

Prior to Space Station, NASA had several advanced studies of space platforms underway. The initial Space Station concepts included one or more polar platforms to meet the needs of Earth observation from this low-altitude Earth orbit. A constellation of co-orbiting and polar orbiting space platforms was included in the Space Station program for several years. Polar orbiting spacecraft are ideal for Earth observations requiring repetitive coverage, so these ideas were eventually combined. Thus, large polar orbiting observatories were used to provide Earth science data.

These concepts coincided with an increasing interest in Earth studies, motivated in part by concern about the human impact on Earth and the sustainable development prospects for the planet. This type of study would require a large interdisciplinary effort to understand the land, ocean, and atmosphere systems of Earth and their interactions.

Other external pressures were also in play. The National Oceanic and Atmospheric Administration (NOAA) experienced budget pressures on the successful Landsat program, which suggested cooperation between NASA and NOAA in collecting and using Earth science data[2]. There were coordinated national programs, such as the Global Change Research Program (GCRP), in which the U.S. was a major participant and NASA a major contributor. In addition, there was also a general expansionist environment in NASA that included an ambitious Space Station and the space science Great Observatories.

[2] NOAA could not get money for Landsat in the late 1970s to mid 1980s because there was a government push to commercialize the system.

These pressures contributed to the establishment of EOS as a large program using few large platforms and many instruments to produce 15-year data sets. The need for a 15-year data set was driven by the desire to include one complete solar cycle. Because the EOSDIS would collect, process, distribute, and archive the data from the flight system, the system was seen as the glue holding the EOS program together.

During this time period pilot data systems for the various Earth science disciplines (e.g. the Pilot Land Data System (PLDS) and the Pilot Ocean Data System (PODS)) were being developed and operated. These systems were stovepiped for the various disciplines and did not span Earth science as a whole. However, the PLDS and PODS spurred discussion of distributed data systems, and became forerunners of the EOSDIS.

The EOSDIS program structure was influenced by concerns about preceding and contemporary missions. Up to that time, most of the data collected on flight missions was processed and held by the Principal Investigators (PI) who built the instruments before the processed data was made available to other scientists. In some cases, especially in the Nimbus program and space science missions, there was a tendency to hoard the data, write the papers, and delay sending the data to the National Space Science Data Center (NSSDC). In part, this was the science community's response to the difficulties they encountered in trying to submit new data products to the NSSDC and other national data centers. The National Science Foundation (NSF), along with several other organizations, complained about this. The process was unacceptable for a large interdisciplinary program.

At the same time, the Upper Atmosphere Research Program (UARS) was in development, featuring a Central Data Handling Facility (CDHF). UARS data went right to the CDHF, where it was processed and made available to all PIs after an agreed time. This model was better one for the EOSDIS, which would be at least an order of magnitude larger than the CDHF that supported only one spacecraft. Another objective was that the EOSDIS data be publicly available to all researchers, rather than only EOS investigators. This could only be done with a large, centralized system.

Another factor influencing the EOSDIS program structure was based on AA observations. They saw that when flight programs got into trouble they would often "borrow" money from the ground and data processing systems, leaving the data processing function short of funds.

The users also determined that the ground system could not be ignored until after launch, as had been the case in the past. The solution was to separate the systems' Unique Project Numbers (UPN) so the flight program manager could not access funds meant for the ground system. The EOSDIS was even put into a separate Budget Line Item (BLI), which meant that NASA had to obtain Congressional approval to add or remove funds. These changes eventually made the EOSDIS more visible to critics and more accessible to Congressional influence and control.

The early vision of the EOSDIS was that it would handle large amounts of data (a petabyte per year) in a centralized facility, and distribute it to a large community. (At one point, this community included 550 PIs and co-Investigators (co-Is)). Despite this large number, one of the conceivers still said, “There would be more discoveries than discoverers.”) This was necessary because of the interdisciplinary work required to study land, oceans, and atmospheres.

The 1986 EOS Data Panel report (the “black book”) listed over 100 top-level requirements and said, “Thus, for EOS to be successful, all of the requirements and systems’ attributes delineated within this report… must be provided.” The data would be available to all at the same time, and the EOSDIS would be all things to all scientists.

The science communities bought into this concept. As someone would say later, “There was broad consensus and broad naiveté.” This attitude, similar to the one carried by those involved with the contemporary Space Station, reinforced that a large system would be built right away. One scientist cautioned that they should first get a system working with existing Earth science data, perhaps using some of the pilot data systems built in the mid 80’s. The NASA AA made this a requirement, and EOSDIS Version Zero (V0) was developed from 1990-1994. Existing Earth science data would later be folded into the EOSDIS, but it did not start out that way.

**START UP**

Beginning studies for the EOSDIS started at Headquarters in the early 80’s. These studies continued into the mid 80’s at Headquarters, at the Goddard Space Flight Center (GSFC), and at the Jet Propulsion Laboratory (JPL), where a concept study was conducted. Various science study panels and working groups were formed, and the EOS PIs were selected in 1988. They were to deliver their algorithms to the EOSDIS in the form of working software.

Phase A studies (~$3M/year) continued until the phase B studies were funded in FY 1989 and 1990[3]. The two study contractors were TRW and Hughes Aircraft. They ran studies out of Goddard’s old Mission Operations and Data Systems Directorate (MO&DSD) where data systems were done, not out of the Flight Projects Directorate (FPD) where projects were done. The DAACs didn’t exist in the phase B studies. Only one DAAC was originally envisioned, and it was to be at GSFC. A second DAAC was quickly added at the Earth Resources Observing System (EROS) Data Center (EDC) in order to include Landsat land data[4]. Several others were eventually added to specialize in various types of

[3] Some scientists complained that they did not have ample input in the phase B studies, despite the fact that there was ample science participation.
[4] The EDC is a U.S. Geological Survey facility.

Earth science data. By 1990, the EOSDIS Project was started in the Flight Projects Directorate, and it received a new start in FY1991.

Some have questioned the use of aerospace contractors for a large data system. In retrospect, a large, non-aerospace, data system contractor (like Microsoft) may have been a better choice. This was, however, a purposeful decision; the aerospace contractors are used to doing big jobs (including data systems) for the DoD and could handle a job of this magnitude. There was also the perception that the same problem had been solved in the classified world, and therefore the defense contractors would be able to do it here. (Later one scientist said, “They don’t know how to do it any better than we do.”)

In addition, the contractors were expected to add the systems engineering talent not available or affordable within the civil service. The Agency had a lack of expertise with large data systems, specifically in systems engineering. Though NASA had been doing complex flight systems for years, it did not have people experienced with large systems. Thus, one of the reasons for consolidating into a large central system was to create a large dollar value to attract an aerospace contractor.

The role of technology in the development of EOSDIS was recognized as important, but the issue was never fully resolved. Debates sprung up about whether with the system should utilize the new, emerging technologies, or whether it should begin with the more mature technologies and futuristically integrate the new ones. It was acknowledged that the technology would not remain static for a 10-year program, but there was disagreement about the evolutionary rate. Some saw disk storage as an issue; others disagreed on the basis that technology would evolve past that need. Still others saw processing capability as more of a challenge, because of the number of products and requirements for multiple reprocessing.

Predicting technological capability in the future was a problem. Some members of the science community thought the Project’s view was too conservative. Eventually, the decision was made to start with existing technologies and to refresh them as new technologies became available. The issue of technological evolution continued to be a major concern and contributed to instability in the system requirements.

> In hindsight, there was a mismatch between the technology lifecycle and the EOSDIS build cycle. The EOSDIS releases were based on a 2-3 year cycle, which was the length of time needed to collect and implement the requirements for a large, centralized system. However, other information system technology was turning over every 6-9 months.

Research also suggested that as a program’s percentage of new technologies approached 25%, there would be a high risk of overrun. Even when the latest technology was incorporated using commercial products, inconsistencies remained. This was because the standard versions across company products caused integration problems in the open distributed architecture of the EOSDIS. These were not good signs for the project.

Requirements for the EOSDIS were generated during study phases A and B, which included requirements recommended by the study contractors in consultation with the science community. Some of these requirements were very detailed and consistent with the pending procurement. However, in retrospect it became clear that there was a great deal of misunderstanding on the approach and disagreement about the development philosophy. The science community saw this as a long term evolutionary development and assumed that their requirements would evolve as the system and technology evolved.

> The instruments were not yet designed, so they couldn't create the detailed design requirements needed for a centralized ECS. There was a disconnect in the requirements process between the science community and the project, and the science community and the contractor.

Aware of NASA's lack of experience in this area, members of the science community did not have confidence that NASA could handle such a large, evolutionary data system. They thought the requirements should be more flexible since the system was supposed to evolve. The EOSDIS Project, however, saw it necessary to have a complete set of requirements in order to specify a system that could be built under contract. The scientists understood the science requirements, they understood the data rates, and they understood how to do algorithms. However, they didn't understand the functional requirements that constituted about 95% of all the EOSDIS requirements (though the small number of science requirements drove the larger number of functional requirements.) The requirements were documented, but the expectations exceeded them. There was never a consensus between the community and the project on these requirements.

Several factors contributed to the instability that plagued the project. One was the three-year pre-acquisition and procurement period that took place during a time of rapid technological change. There was also a procurement blackout that precluded people from discussing the requirements. There was not enough time to work out the requirements (beyond listing them initially) prior to acquisition. More time should have been spent in a less pressurized environment, listening to people who understood the technology. Voices of competent people weren't being heard. Rather than "vertical" discussion, there was only "horizontal" discussion between the experts.

> The different views on the project requirements were described by one team member as two separate cultures: a "box" culture, or constrained, well-defined boundaries, and a "cloud" culture, meaning the boundaries were amorphous.

The box culture, which included most members of management, wanted to build a system according to a set of fixed requirements. For a traditional procurement, it was necessary to have a set of requirements that must be provided to the contractor for the system to be built. The cloud culture, which included most of the scientists, wanted the requirements to change as knowledge was gained. Both views were justified according to the original vision of the EOSDIS, but the challenge was separating what was to be done in the "box" and what was to be done in the “cloud”. Since this issue was not solved early in the program, it continued to plague the development.

In February 1996, a little over two years from the planned launch of the first EOS spacecraft, there was a freeze on the computation and storage requirements. Some interfaces were baselined to preclude further delays. (See the next section on implementation.) The instrument teams were still unsure of their final requirements, but the database had been set in concrete. Some members of the instrument team were angry about the freeze; their opinion was that the requirements should continue to evolve.

Looking back, one contractor observed that this pre-acquisition phase was the major problem. Everyone, including the science community and the government, should know what they want to buy prior to the procurement phase. But because of this different view on fixed versus evolving requirements, that was not the case.

Following the studies, procurement activity lasted several years. Because of the dynamic nature of the program and its changing technology, the blackout period during procurement was detrimental, yet unavoidable. During this time, a new NASA Administrator came onboard. He sent both bidders back to revise their proposals based on cost realism, as he thought they were unrealistically low.

Hughes won the EOSDIS Core System (ECS) phase C/D contract in 1992, and the contract was signed in 1993. The value was about $800M, which was almost double the original bid before the re-proposal. The contract was an end-item specification contract, which meant that the end item ECS was specified by the government and expected to be delivered by the contractor. The problem was that, as discussed above, there was no agreement between the parties-- the government, science community, and contractor--about what the end item was. Despite the existence of the specification, the views on system evolution diverged.

The Administrator also directed that there be a Total System Performance Responsibility (TSPR) clause. This held the contractor responsible for making the system work within budget, even though there was work within the DAACs outside of the contractor’s control. TRW won a separate contract for the EDOS in 1994.

Shortly before the contract was signed, other Agency problems caused cost pressures to build. An external Red Team review of the EOS program was established by the Administrator. ***As a consequence, all the projects in the EOS program, including EOSDIS, were charged with reducing their budgets by 30%. These reductions came with the direction that there was to be minimum impact on science.*** They required some negotiated reductions in the ECS contract requirements and reduced critical flexibility. The impact of these reductions, which was dramatic, is discussed more in the next section on implementation.

Hughes had structured an ECS development approach that included the parallel development of subsystems, coming together for integration of the EOSDIS. On a schedule, this is a "waterfall" of events culminating in the final delivery of the system. This is not a bad model for the development of a familiar system that has been built more than once. It depends on having a stable, well-defined specification. However, many believe that this project was over-compartmentalized, even for a waterfall development. There was little communication between the parallel developments inside the contractor's team.

Some have suggested that this may have been the wrong development model, considering the evolution desired and the rapidly changing technology. A more appropriate model was a "spiral" prototyping model. According to this model, the first cut of the requirements are taken, then the most difficult part if built, tested, and used. If it is satisfactory, then the next part is built. Parts are iterated, built and tested, built and tested again, creating a spiral that completes development. This model was not in frequently used within NASA at the time. Instead, the plan of action for both the EOSDIS and contemporary Space Station Freedom was "do it big, right away." Based on this structured approach, discussions continued and technology rapidly changed, even as the contract was being signed.

Early in 1993, Hughes held an open meeting at GSFC to brief the science community on the proposed EOSDIS architecture. This architecture was based on the requirements from early on in the study and procurement processes. The design was, as stated, highly centralized. It was perceived by some users to be a dead end design, but there was little feedback at the meeting. Groups within the science community were talking, but nothing got back to the project. Better communication at this meeting could have helped potential problems to surface earlier.

The ECS contract with Hughes was signed in March 1993. Later that summer, Hughes conducted an open Systems Requirements Review (SRR). The science community was represented by the EOSDIS Panel (a.k.a. Data Panel[5]) of the EOS Investigator's Working group (IWG), but there was also a broader attendance at this review. Although there was little feedback from the science community at the previous meeting, they had a significant negative reaction at the SRR. They saw a rigid, centralized system being developed under an end-item specification contract, and it was not at all what they anticipated with respect to evolution. They could not see how this design could evolve.

The requirements they had previously submitted were cast in stone, with the assumption that newer hardware would be infused when it became available. In an attempt to work the issue after the SRR, the project held a retreat for the project, Hughes, and NASA Headquarters. Hughes was told to update the architecture and decentralize the design to address the community's concerns. This was to be approved by an external committee of senior scientists and managers by December 15. Some significant architectural changes were made, and there was general agreement at the December meeting. In retrospect, it is

[5] This Data Panel was different from the one that did the 1986 report.

clear that the changes were not enough to assure a smooth development; there was still no change in the basic philosophy that everything would be built simultaneously in parallel.

At the same time, there was a meeting between Hughes and the new NASA Administrator[6]. The motivation for the meeting was unclear. It had been reported that he suspected Hughes could not do the job within budget, and that some of the scientists were complaining directly to him. There was also a lot of budget pressure at the time, especially on the Space Station and EOS, and many in Congress were critical of the EOS program.

> The administrator made it clear that he expected Hughes to do the job according to the contract, or it would be cancelled.

Not surprisingly, he cited the cost realism re-bid proposal that had increased the cost from the original proposal.

There were people at Hughes who found this to be a significant meeting. Hughes' original plan was to bring in experienced talent from the classified world in order to make the necessary changes to the EOSDIS. But after the meeting with the Administrator, they decided to hunker down, listen to his words, and go back to the NASA way--which was following the letter of the contract. The expansionist environment present at NASA years earlier had disappeared. Thus this meeting, while significant, should not be overly emphasized. It was as much a symptom of the environment as of any other factor. It was in this time period that EOS was impacted by significant funding reductions.

But the meeting may have had another impact. There were known weaknesses in the Hughes proposal which had been identified by the Source Evaluation Board (SEB). The hard line taken by the Administrator, along with the desire to have Hughes live with their proposal meant that these weaknesses could not be corrected before the contract started. They were built into the start of the implementation phase.

The Administrator also convened a set of his former colleagues from the classified community to look at EOSDIS. They concluded that the approach was consistent with their best practices at the time. Whatever the motivation for these two events, the SRR and the meeting with the Administrator signaled a rough start for the work ahead.

## IMPLEMENTATION

> Early on in the program (1990), the role of the EOSDIS was expanded to include not only data from the EOS flight program, but data from other Earth science missions such as Tropical Rainfall Measuring Mission (TRMM), Landsat, and UARS. As a consequence, the project changed from being the Earth Observing System DIS (EOSDIS) to the Earth Science DIS (ESDIS)[1].

[6] It should be that noted that the Administrator was involved with the TRW team bidding on EOSDIS before he became NASA Administrator. He, therefore, stayed distant from EOSDIS matters as a consequence.

The change was motivated by the fundamental assumption that all Earth science data from various NASA missions would be used together. There would be new, and more integrated ways to accomplish Earth system science.

Shortly after the new start for EOSDIS in late 1990 (FY1991), it became essential to deliver improvement in data systems early and to have something to show for the money being expended. EOSDIS sponsors, Congress for example, were anxious to get something delivered. This capability, termed Version 0 (V0), included a migration of much of the earlier pilot data system work (e.g., PLDS), and depended on this earlier development work.

There were many motivations for V0: It would be useful whether EOS missions flew or not, one could learn about data system development by managing these heritage datasets and making them better accessible to the community, these data sets could be managed by distributed data centers (which came to be known as DAACs), and the EOSDIS could evolve from this development. The World Wide Web (www) came into being about one year before the completion of V0 and was to some extent incorporated into it. The V0 system was completed by August 1994; it is still operating. The V0 delivery demonstrated that success could be achieved by concentration on a narrowly focused goal.

Even early in the implementation phase, (1993-94) milestones were missed and schedule was lost. By design, the government scheduled many data item deliveries early on. Hughes could not keep up with the demand and was delivering “shell” documents containing little content. This made it necessary for content and delivery renegotiation, which pushed the schedule back even more.

> In 1995, the ECS Preliminary Design Review (PDR) was held without the final set of requirements which is the prerequisite for an end-item specification contract.

In 1993 the ESDIS Project was moved from the Flight Projects Directorate (FPD) to the Mission Operations and Data Systems Directorate (MO&DSD) where the original studies had been conducted. The MO&DSD, since reorganized out of existence at Goddard, was the traditional location for ground and data system development. It was institutionally structured, not project structured, meaning that the systems they developed supported the institution.[8] There was no project management experience in the MO&DSD, so the project was run by an Associate Director, formerly of the FPD, who had prior flight project management experience. This move symbolized the ambivalence of GSFC management for locating and running the project. (The project would eventually be moved again in1997, this time into the Flight Projects Directorate (FPD).) EOSDIS is

[8] In fact, while the project was in the MO&DSD, there was pressure on the Project Manager to make sure he supported the institutional developments.

primarily a data system, an area of expertise found in the MO&DSD, but it is also a project, and that management expertise lies in the FPD.

The Center did not have expertise in the project management of large data systems, although it had significant capability in flight project management. One person associated with a competing contractor observed that the government people were more actively participating than those they expected. This was a sign of the inexperience. In fact, there had been six project managers over a period of 12 years, and most of them did not have the basic level of experienced deemed necessary for this particular job in the private sector. Despite that, all of them contributed under difficult circumstances, and the later ones were instrumental in completing the system.

This lack of experience also existed at the program level at Headquarters. Although they possessed the vision to conceive the EOSDIS, they didn't have the appropriate experience to manage large data systems. Shortly after the new start of the EOSDIS, the project was granted limited authority to recruit a new EOSDIS Program Manager. A senior position was provided for this function in 1994, but by then it was too late.

This situation continued even after program management was moved to the Centers in 1996. Before this move, there were about 10-12 people experienced in data systems (although not necessarily program management) at Headquarters. With the program move, a number of these people retired or changed jobs. Only a few made the move to the Center for the data system program management, and they were not experienced program managers.

This lack of experienced leadership was most acute at the top. By the mid 90s, all the original prime movers and leaders at Headquarters had moved on. There was no program champion at Headquarters who could set major policy, decide between competing agendas, and who could direct the community <u>and</u> the project. There was no longer a division-level organization focused on data systems and issues. As a result, the project had to react to competing recommendations from various reviews without the executive leadership from Headquarters to decide between them.

The issue was manifested early as a lack of change leadership following the contentious SRR. It was evident in the inexperienced government team, in the lack of architects and leaders in the contractor team, and in the lack of a unified science community leadership to bring the project team together.

> There was a clamor by the science community to design an evolving system. The problem was that a rigid, performance-based contract was already in place.

Change, though possible, had to be managed and backed by additional funds for change orders. The project recognized the complaints of the users and acceded to the requested changes. This was done informally; changes would be suggested in open meetings and accepted. Open communication is necessary between users and developers for successful development, but it cannot be used to direct contract changes. Too many people other than the Contracting Officer were giving direction to the contractor. The users would talk directly to the contractors. Some scientists were making requests to change the requirements even though they had no responsibility for implementing those requests[9].

Neither the government nor the contractor had an adequate configuration control process. Change requests would not be acted on, and some remained open for years. The inexperienced project people controlled neither the expectations of the users, nor the process for changes in the contract. On top of this, the contractor acceded to any and all requests. The whole process resulted in delays and cost increases, and as a result, none of the principals (the project, the contractor, the users) got what they wanted. The project did attempt to make changes, as shown by the architectural changes made after the SRR retreat. (After which a complaint was made by the scientists that the project didn't adopt any of the suggestions.)

They also facilitated work at the contractor's plant using project personnel. The changes were not significant in the context of the existing contract and architecture. Changes to the basic approach would have required a "change management approach." Even at this late date, with the differing expectations of the parties involved, a very experienced management team should have been able to manage the conflict, regroup, identify doable objectives, and start to move toward these objectives. This did not happen.

The ECS contractor's actions (Hughes) contributed to the difficulties. Some suggested that Hughes removed a good proposal team after the win. To replace them they hired inexperienced programmers off the street instead of providing a top-flight development team for this very complex project. This is consistent with the comments of a former Hughes employee, who said they refrained from hiring top-flight talent from the classified world after the meeting with the Administrator, partly out of fear of overrunning the contract value.

The Hughes staffing approach was also criticized. They staffed very quickly with inexperienced people, rather than waiting to find more experienced people that they could hope to retain. Mid-level, experienced managers were in short supply. As it was, the turnover rate was very high--approaching 38% at times, and rising even higher in the mid-management ranks[10]. Efficiency suffered dramatically. At the nadir, the production rate was 3.28 lines of code (LOC)/day/person. This compares poorly with a more normal rate of 15LOC/day/person or even higher for motivated or talented programmers.

Another former Hughes manager acknowledged that Hughes hadn't stepped up to the job. After the contract was awarded, Hughes brought in a project manager who had successfully managed a large information system development for another government

[9] Some suggestions could have been made as simple suggestions, but they were interpreted as direction. If this happened, the project should have controlled the process.

[10] This was exacerbated by the dot com boom in the mid 90s, a time when any breathing programmer could be offered a lucrative job at any time.

agency. However, he was reassigned and replaced by a project manager whose experience was neither project management nor information systems.

The external environment also changed from the early to mid 90s in reference to budget and political issues. EOS was originally meant to be a $30B program at completion, or $17B through 2000. The EOSDIS had about $6B of this total. In the beginning there was also a lot of contingency money for the EOSDIS. However, at this time the Agency was struggling with budget cuts and Space Station overruns.

The political climate changed as well. There was now less support for EOS in Congress. In the early 90s, the National Space Council made a recommendation to cut EOS due to its large size. One of their suggestions was to reduce the optics size to reduce cost. This would have required a lower altitude and was not done.

The combination of these factors, project difficulties, and external pressure, led to a series of reviews and re-baselines for the EOS program. The flight program was changed, causing it to go to smaller spacecraft and smaller launch vehicles. The EOSDIS was also impacted. The $17B meant for EOS through 2000 was eventually reduced in a series of steps, first $11B, then $8B, and finally $7B by mid-decade. The 1995 EOS Execution Phase Project Plan (EPPP) summarized the re-baselined EOS program at $6.1B, with the EOSDIS receiving $1.9B. (The ECS contract is the single largest part of the EOSDIS budget.)

> But as the budget was reduced, the requirements were not changed accordingly. The original idea was invoked: if the money available for refresher costs could eventually buy better capability, then existing capability can be maintained with less money.

The necessary functions for collecting, processing, distributing, and archiving data remained, but others were impacted. First the reprocessing capacity was reduced, and then user services were impacted. Significant reductions were made to contingency funds, because some functions originally planned as part of NASA’s institutional budget (e.g., EDOS and network capabilities) were now EOSDIS responsibility. There was concern that the EOSDIS would be cancelled. One report said a member of the prestigious external review teams had threatened exactly that. This contributed to the project’s desire to please the community by accepting the proposed changes.

The world was changing rapidly in other respects besides budget and politics. The World Wide Web (www) was starting to come into its own. Some technology was moving faster than anticipated, while network capacity was not. Scientists facing the difficulties of interfacing with the EOSDIS, looked at the Web and increased pressure for a more decentralized EOSDIS approach. Some scientists liked the fact that they could use the

Web to distribute data after it was processed by scientists with the necessary expertise. The web was suggested as a way to distribute this data[11].

A number of alternative, more decentralized architectures were studied and proposed by the science community. They were concerned about money because they thought the ECS was "eating their lunch," which was not entirely accurate. Not coincidently, the proposed system would also distribute the funding to the users and put control into the hands of the science community. Even though there were many self-inflicted problems in the ECS, the external budget reductions were having a significant impact. In particular, the enormous decrease in available contingency funds, and the inability to accurately estimate the cost of user services, undermined the project's ability to manage the changing environment and the opportunities it presented.

Later attempts were made to reduce the large workforce of the ECS contractors (which numbered 720 full-time-equivalents (FTE) at the highest point) and bring doable work in line with the decreasing budget. An ECS contractor presented the impact of reducing the large workforce to a Maryland Senator. Resulting pressure on the Agency forced them not to make any significant changes to the workforce, which increased the overall budget pressure. Ironically, the original idea of using a large, central system with a large dollar value (a plan adopted to assure the attention of a large aerospace contractor) now made it difficult to avoid the political spotlight.

When a high-profile, high budget government program is in trouble, external reviews are unavoidable. The EOSDIS was no different. Numerous independent review teams, including some from the National Research Council, descended on the project. One of their major objectives was to ascertain the cost of individual services. The idea was then to reshape the program by determining which services could be dropped and how much could be saved. This goal was never met. One of the reasons, though certainly not the only one, was the consolidated architecture of the ECS. There was never a clear delineation of the cost per function. Everything was connected to the infrastructure of a monolithic system. Without knowing the cost per function, it was impossible to reduce cost. The problem was exacerbated by the big-build approach and EOSDIS was never restructured based on the recommendations of the external review teams[12].

By the mid 90s, the EOSDIS was behind schedule and over budget. Many doubted that anything would be delivered. In 1996, despite regular reviews, it became obvious that

---

[11] Although the web works well in a text environment, it is not obvious that the sophisticated data of EOS would be as well-handled. One interesting observation was that the web protocol was market-driven; it came from below. It was the opposite on the DIS.

[12] However, as a result of one review, an initiative was started that recognized the importance of decentralization and innovation. The review (NRC, LaJolla CA, 1995) pushed for a Federation of Earth Science Information Partners (ESIPS) which would conduct experiments in research and applications in data systems. The emphasis was to be on innovation in science and technology. For approximately $12M to $15M, 24 Cooperative Agreement Notices were initiated by various groups to start this process. This Federation is self-governing and has expanded to include the DAACs. It is a continuing operation.

ECS Release A, which was destined to support TRMM, would not be ready on time.

The project cancelled Release A and used V0 and back-up capabilities developed at the Goddard and Langley DAACs to support TRMM. Too much had been promised to the large, diverse science community. While it was easy to reach agreement among developers of a small mission, the trade space in the EOSDIS had become unmanageable.

There were too many common denominators being sought. There was a mismatch of expectations. Many in the science community saw it as an R&D project, while the government managers just wanted to deliver something.

**RECOVERY**

During the course of the program, there were several management changes both at the EOSDIS Project and at Hughes. The key changes occurred at the same time in late 1996, and then again in 1998.

In 1993, the NASA Administrator became concerned about the EOSDIS budget and asked a pre-NASA colleague to review the EOSDIS. This person, a government employee in the classified world, concluded that the EOSDIS was destined to fail for several reasons. One, the program was too ambitious for the money allocated. (By this time the EOS budget had started to come down.) Two, there was too much contention built in. He felt the scientists had too forceful a voice, and that Congressional support would be lost if the government didn't respond.

The review caused destructive tension. The reviewer left the government in October 1996 and joined Hughes as a Vice President (VP). The GSFC Deputy Director then requested he run the program from the contractor side.

Another, more experienced project manager with several successes to his credit took over the ESDIS Project in October of 1996. The realization set in that the EOSDIS would have to be restructured and more realistic goals set. They started to tighten the requirements by adding more discipline to the change process, identifying more doable objectives, and adopting a more orderly, incremental system.

A goal was set that "by September 1997 we will demonstrate this capability." That goal was then met. They realized that everything originally promised could not be delivered on the basis of current performance and ongoing budget reductions. There was an attempt to define what could be delivered. They addressed the excessive workforce level on the ECS contract. In essence, they attempted to get more management control of the project at the expense of some original goals. The ability to achieve these goals had already been eroded by previous management and a reduced budget.

The new Hughes VP managing the ECS contract determined that the program needed a strong systems management. (Systems management is the combination of project management and systems engineering.) He thought they had a short time to turn the program around before it would be cancelled. He also saw a major deficiency in the systems engineering capability of the Hughes development team. This same observation had been made years earlier by members of the government, but the problem was never corrected. He hired two colleagues as systems engineers (who were also former government employees from the classified world) shortly after he came on. These new systems engineers were experienced with large data systems. He made several other key hires as well to strengthen the core of the team. The new people hired to lead understood systems engineering for large systems[13]. A good working relationship was established between the mid-level managers in ESDIS and these systems engineers.

More control and project discipline was applied, more doable goals were set and achieved, and a new project management mentality was imposed, but there were still problems during 1997-1998. First, Hughes kept trying to deliver modules, but they had difficulty meeting schedules. Second, there was still concern about the excessive workforce. Third, the government/contractor relationship, although more stable and realistic, was still confrontational. Structure and limits had been imposed, and progress had been made, but both teams understood the need for change. The bottom had been reached, and the basis for a turnaround was in place.

Raytheon purchased Hughes in December of 1997. The VP who brought in the new team left shortly thereafter, but the key people remained. The general opinion among the government people and the remaining Hughes employees was that Raytheon was more concerned about effective, responsive project management. One example was the establishment of regular communication between a Raytheon VP and the GSFC Deputy Director.

As the government started to get control of the situation (1997-1998), the Flight Operation System (FOS) failed a major test that uncovered a serious problem. The FOS was to be the operating system for all EOS spacecraft. It was prepared first for the AM-1 (later renamed Terra) spacecraft, which was originally scheduled for a 1998 launch[14]. However, during the press of business that surrounded the SDPS, the FOS had not been watched as closely as it should have been. A full-court press was then applied by the government and the new Raytheon team. After considering a number of options, the existing system was dropped entirely, and they turned to a new, commercial-based system.

[13] Ironically, the core in place was a team of Hughes systems engineers with experience on large data systems in the classified world. This was said to be the early intention of Hughes, but it was reversed after the meeting between Hughes upper management and the NASA Administrator. Following that meeting, Hughes followed the letter of the contract. They hired a less-qualified team to keep from overrunning. An original goal of the conceivers was to attract a large aerospace firm capable of bringing significant systems engineering talent to the project.
[14] The FOS was subcontracted to Lockheed Martin by Hughes.

This system was assigned to Raytheon/Denver. Within a year, the FOS was completed using the new system. The FOS delay also contributed to a launch delay for AM-1, but the government/Raytheon team managed a quick recovery. The new version of the FOS was termed the EMOS-EOS Mission Operations System.

In early 1998 a new Office of Earth Science Associate Administrator (AA) was appointed. He had a better understanding of the EOSDIS than most of his predecessors. His experience included development of the PLDS, and his idea for EOSDIS recovery was to use the philosophy that resulted in that project's success.

The model was a concentric one. At the core were the tools needed by anyone using the EOSDIS, such as search tools, formats, and retrieval systems. The next outer concentric circle would be made up of functions common to two or more groups of users. The outside ring would hold very specific items needed by only a few. The idea of focusing on a small, doable core was the same one being worked by the ESDIS Project at the time. Later in 1998, a new ESDIS Project Manager who had more data system experience was assigned. She was instrumental in carrying this core idea into implementation. Her leadership was a key factor in achieving the cooperation necessary between the EOSDIS project team and the new contractor Raytheon.

A 1998 large budget over-guide submitted for the EOSDIS by GSFC, triggered some major changes by the new AA. Headquarters directed that the Project do a design-to-cost EOSDIS. This was necessary due to a constrained budget, and also the need to produce an EOSDIS before they lost all Congressional and community support. These directions coincided with another external review and a critical report by the Littles' Team. While the Littles' review was just one of many done on the EOSDIS, it came at a time when the Agency was considering organizational change.

Shortly after the Littles' review, there was a program reorganization. This reorganization changed the GSFC program management structure to resemble the structure outlined by the NASA Procedure and Guideline (NPG) 7120.5 entitled NASA Program and Project Management Processes and Requirements. The EOSDIS was then put under a program office in the Flight Programs and Projects Directorate (FPPD) at Goddard (formerly the Flight Projects Directorate (FPD)).

The event that spurred the turnaround was the 1998 development of options by the ESDIS Project, specifically Option A+. At the time, this was the large budget over-guide, and the new AA made the direction to build-to-cost. Option A+ was the first time all agreed that requirements should be put in a "box" (described previously in this report).

There was a more disciplined approach to completing fewer requirements. Science requirements were limited and prioritized, with only 17 high-level requirements. These 17 requirements were restatements of requirements at a high level, but they were stated in a prioritized order that was vetted with the science community. They were worked on according to this order and were stated in terms of functional capabilities to be tested and verified. Thus the ECS SDPS became an incremental development with frequent deliveries of completed capabilities.. ***Option A+ was to form the core of the revised ECS SDPS using the new model. It created stability for the project, convinced Congress that something would be built, and signaled to the science community that all original promises would not be fulfilled.***

Option A+ included restructuring the entire ECS contract with Raytheon. Some of the data processing was removed from the contract and assigned to Science Investigator-led Processing Systems (SIPS). These science-led facilities are essentially decentralized processing facilities at the science institutions. The SIPS scientists do their own requirements, algorithms, software, and processing for the instrument data in the SIPS. They are funded separately for those tasks from the ECS SDPS. The data is received from the ECS SDPS, processed, and returned to the ECS SDPS for distribution and archiving. The SIPS represented the decentralization of data processing that had been recommended in lieu of a central facility. They represented the next layer out from the core in the concentric model.

> The SIPS operations were simpler and more austere than the services envisioned in the original ECS SDPS concept. This is due to the self-generated requirements and the funding received. The people who generated the requirements were the same ones responsible for keeping SIPS on budget. This was in contrast to the original ECS SDPS model in which the requirements made by the scientists involved did not come out of their direct funding. In the latter case, there was little incentive to restrict requirements.

In fairness to the science community, they were originally encouraged to believe that the ECS SDPS would be everything to everybody. (They were also concerned about costs and frustrated when their requests for cost accounting were not received.) In addition to the transfer of data processing to SIPS, some functionality was removed from the ECS SDPS. This included several user services. The result was that the core contract work was reduced and became more manageable. There was also a better relationship between the requirements and the funding in the SIPS.

The Raytheon Program Manager realized that the ECS SDPS would have to be reduced to make it more “doable.” He agreed to a reduction, and achieved the desired results. The loss, however, was a marked reduction in the services originally envisioned for the users.

Also at this time, cooperation between the government and contractor teams became significantly improved. Contact was improved, and a partnership developed. The idea that “we’re all in this together” superceded a more confrontational approach. There was more support from Headquarters; the AA, understanding the nature of the EOSDIS, now supported the tough choices that needed to be made. He helped protect the project from additional outside reviews that had previously taken a considerable toll on the project.

Option A+ worked. The original requirements of the EOSDIS had not been met, but an acceptable match had been made between the necessary functions and the resources available. This, of course, is the goal of any project. The EOSDIS implemented at the EDC DAAC supported the initial operations for Landsat 7, subsequent to its April 1998 launch. Additionally, the EOSDIS with the restructured FOS (EMOS) supported the EOS AM-1(Terra) launch in December of 1999.

EOSDIS continues to support EOS and other Earth science missions along with the SIPS in collecting, processing, distributing, and archiving data, as well as operating the flight missions. It also works with other organizations related to the EOSDIS in the previously discussed “outer concentric rings.” These include the Federation of Earth Science Information Partners (ESIP, see footnote 10), and the Regional Earth Science Application Centers (RESAC).

The ESIP federation brings government agencies, universities, non-profit organizations, and businesses together to make Earth science information available to a broader community. It started in 1997 through Cooperative Agreement Notices (CAN) and uses a very deliberate and incremental approach. Type I ESIPs concentrate on producing standard products, on a strict schedule, in the highly reliable DAAC environment. Type II ESIPs are responsible for data and information products and services in support of Earth system science. These products and services are developmental in nature and are used when an emphasis on flexibility and creativity is essential to meeting advanced research needs. Type III ESIPs are expected to provide information and services beyond the Earth science research community. (http://esipfed.org)

NASA selected nine geographically distributed, academic/government/industry consortia in 1998 to form seven RESACs. Each received a grant to apply NASA’s Earth science research to well-defined problems of local interest. Some of these included precision farm management, landcover/use mapping, urban sprawl, and fire hazard management. (http://www.esad.ssc.nasa.gov/resac)

## LESSONS LEARNED

The ESDIS development from the mid 80s through the 90s occurred during a period of dynamic change within the Agency. As a result, the project felt the effect of numerous external pressures throughout development. A number of internal decisions also led to delays, cost overruns, and animosity between project teams. In retrospect, many lessons can be drawn from this development, both in program and project management. Such lessons can be applied to ensure smoother development of future flight and ground systems.

Many things that should have occurred at earlier stages were not obvious until much later; the lessons were learned after the fact. This is the reason the lessons learned are being covered after a discussion of the full development story, rather than being interspersed throughout the appropriate sections. Not unexpectedly, many of the key lessons learned were from the start-up phase. Had some things been done differently, many of the later problems could have been avoided.

Early Concepts

1. Don't overreact, or (to mix metaphors) fight the last war or let the pendulum swing too far in the other direction.

One of the problems of the past was data hoarding and distribution delays after launch. This problem could have been solved with a modest approach that simply addressed these issues—without a large, centralized system promising everything to everybody. Although there were other motivations for the large centralized system (such as the need for interoperability of large data sets) the reaction to past problems largely drove the design.

This was also true of the separation of the EOSDIS UPN and BLI from the rest of EOS. While it seemed helpful at first, it provided extra visibility for the EOSDIS, and therefore, more vulnerability. The solution was more extreme than necessary to address the original problem. The size of the ECS could have been significantly reduced—and science community support for processing and reprocessing could have been guaranteed from the start—by relying on science teams to produce their standard data products instead of delivering algorithms to a central facility. The science community had progressed in its expectations and recognition of data sharing, now seeing it as essential.

Start-up

2. Know what you want to build and be able to define it.

Although there was an ambitious vision that the EOSDIS would serve all users from science researchers to K-12 students, there was never an overarching concept for implementation. The requirements were known, but they were too diverse and all-encompassing. There was tension between the mission critical requirements (e.g., spacecraft command and control, and acquiring and processing data to an acceptable level), and the scientific understanding of progress and the evolution of technologies. The latter would drive the reprocessing and reanalysis of data.

There was never a consensus among the stakeholders regarding fixed versus evolving requirements. When the rigid end-item spec contract was developed, it was not matched against an overarching system concept. Attempting to do so would've sparked the realization that the contract concept didn't match the system concept, and that a smaller contract with more flexibility to develop outwards was more practical. System definition would not require a rigid end-item spec contract covering the whole system at once. Rather, it means defining the development process in a logical manner following the overall concept.

The following three lessons are related, and are thus labeled 3a, 3b, and 3c.

3a. Acquisition strategy must be tailored for any system where the user needs are difficult to articulate and subject to technological evolution and enhancement.

Start in the pre-acquisition phase with the definition of the conceptual system architecture based on broad system objectives. (See lesson 2.) Tailor the acquisition strategy and associated contract in such a way that development is evolutionary (i.e., comprised of multiple blocks/releases where each block contains multiple increments.) This incremental development allows the users to get smarter with time by understanding what works and what doesn't.

Incremental development allows demonstration of functionality at each stage. This was the most universal lesson learned according to the principals at Headquarters, GSFC, within the science community, and among contractors. This means following a build-test/ build-test plan in which functionality is demonstrated at each step. The final step is putting the last piece in place—not connecting several separately developed subsystems and hoping they work together. This allows corrections to be made in smaller increments, and if necessary, a change in direction. It also allows for an interim operating capability.

This was achieved with the ESDIS Project developed Version 0 (V0) operating system in 1994. The incremental development approach facilitated the needed interaction between the science users and the developers. Changes were easier to make at these smaller increments.

3b. A build-it-by-the-yard approach is desirable to maintain cost control while allowing flexibility for evolutionary changes--especially for large, complex systems expected to evolve during an extended development period.

"Build-by-the-yard" implies starting system development with a small core common to all users, and then incrementally adding functions for more specialized users. If an end-item spec contract is used, it should be as small as possible to maintain development flexibility. Room should be allowed outside the core to accommodate the more unique and changing requirements in the outer concentric circles. The development proceeds concentrically outward as more is learned during the incremental development described in Lesson 3a.

Not only is the development done incrementally, but decisions based on the possibilities within the overall system concept and resources can be made according to progress. If the technology does not progress, or core functions are more costly than expected, functions originally planned for the outer concentric developments can be cut. One must be able to react to new technology (especially information technology (IT)). Procurement of the EOSDIS more closely resembled an IT service than a product. This service is always harder to specify all at once in a concrete way.

This buy-it by-the-yard approach allows for evolution over time and is opposite to the waterfall or turnkey approaches that were attempted. A turnkey or waterfall development is appropriate when the product is well-envisioned, like an F-16 or a series of NOAA spacecraft that are copies of an original. It even works when the specifications are well-defined for a science mission spacecraft. However, for a dramatically new development like the EOSDIS, there are many unknowns. A waterfall development assumes that everything is known; solid specifications go into the front end and then a product emerges.

The development of Option A+ represents the application of this lesson[15]. Option A+ became the new core. It did not include everything originally promised, but it did possess the necessary for the system to be built. Other requirements not included in the new core were done in other ways, through SIPS for example.

3c. Flexible options must be available for the outer concentric developments

Once the small core development is assured, more innovative methods can be used to develop the outer concentric circles. This could take many different forms. For example, an integration contractor could be used, perhaps the same one that developed the core capability, instead of one large contract responsible for everything. The integration contractor could then use several smaller developments for the parts of the system in need of integration. These smaller developments could be tested simultaneously, and the unsuccessful ones could be discarded. The integration contractor may have been Hughes or Raytheon, but the others would be small entrepreneurial companies or academic institutions with software expertise. This is equivalent to a market-driven approach; try some approaches, and discard Those that are unsuccessful.

With a small core in place, the outer circle developments could benefit from a spiral development. This type of development involves building and testing. If the test is successful, you move on to the next step. Of course, there is no guarantee that all steps work; the approach would not have been problem free. It is similar to the entrepreneurial approach mentioned above. However, the approach would have given ample opportunity to learn and change direction without any risk to the mission critical components. It may have seemed risky to managers who were used to end-item spec contracts, and would require strong leadership.

> It seems clear that the acquisition strategy (as envisioned in the early stages to attract the right contractor) was not appropriate for this large, new, rapidly changing IT development.

The development of the SIPS does represent a positive application of this lesson. The SIPS are on the next outer development circle from the core; they are not part of the core. The separation of the SIPS work allowed Raytheon to concentrate on the much smaller core work. The SIPS involve entrepreneurial scientists with narrower interests than needed for the EOSDIS. The persons involved were responsible for both the

[15] More correctly, Option A+ was a priority based set of capabilities that resulted in an incremental development approach. The approach was budget driven and designed-to-cost, however, it represented a new core.

requirements, and the budget to accomplish the tasks. This approach reduced the contention on requirements, allowing them to stabilize. It also assured the scientific experts' direct involvement in the design and development of data processing algorithms. This eliminated the contractor's duplication of this expertise. Had an alternative acquisition strategy been the basis of the project from the outset, the project implementation shortcomings could have been avoided, and the value of money spent could have been improved.

4. <u>Control expectations; tell the truth about capabilities.</u>

A project is designed to produce a certain product within a certain time and resource allocation. The time and resources are not unlimited, so the requirements on the product are not unlimited. During the expansionist era of the late 80s, the EOSDIS focused on one large end product and expected it to be done in one large step. It promised all things to all people—promises that were not sustainable even if the original budget had remained in place. When the budget was reduced, there was no hope of achieving the original promises.

A critical responsibility of program and project management is to keep everyone on the same page in reference to the available resources. This is a major job under stable conditions, let alone dynamic ones. Expectations were not controlled on EOSDIS until the mid to late 90s, even though the budget began dropping years earlier. Nor had they been controlled while the project itself was making the decisions that led to delays and overruns. Constrained resources will always be a challenge for a project. But even if there isn't enough money to do it all, it's still possible to practice good science with the available resources.

5. <u>Choose the appropriate organizational structure, staff it accordingly, and stay with it.</u>

The development of the EOSDIS is a complex, but typical, project-type activity. It involves the management of technical requirements, schedules, and financial resources. In this sense, the project correctly belongs in (what is now) the Flight Programs and Projects Directorate (FPPD). However, starting the early studies in the MO&DSD was appropriate for a data system. Because of the technical expertise in the MO&DSD, a strong partnership between the project and that team was necessary. But the mid-project move from the FPD to the MO&DSD (before being moved back to the FPD in 1997) was a mistake which was later acknowledged.

The reason for the move was the Center's inability to find the right blend of data system talent and project management experience. This was a problem across the Agency. However, the answer was not to move the project back and forth. While the project was in the MO&DSD, there was pressure on the project manager to make sure he supported their other institutional developments.

Strong, experienced systems engineering—and the associated processes for both the government project office and the contractor team—are essential ingredients for success.

Similarly, the contracting officers must also be strong and experienced, specifically on a large program where not attending to details can result in magnified consequences.

6. <u>Keep the flight operating system (FOS) tied to the flight segment.</u>

An original motivation for consolidating the ground system to include the FOS, as well as data capture, processing, distribution, and archiving, was to protect its resources. This protection was against raids by the flight segment when they ran into overrun problems, which happened often in the past. Another motivation was the necessity of operating multiple missions from one control center.

Consolidating the FOS with the ground system seemed like a good idea at the time, but it caused problems later. The FOS development is closely linked to the development of flight segment, specifically operations, integration and tests, procedures, software, and commanding. Although data capture and processing functions can be separated from the flight segment, the same is not true for the FOS. It is developed by the flight segment team, even if it is later turned over to another contractor (along with other similarly developed missions) to be operated in a common mission control center[16]. In this case, the FOS was separate from the flight segment development, but was tied very closely to the development of the SDPS. This was unnecessary. In addition, proven FOS development methodologies were not used; it was built with too few development versions.

As a result, development problems were hidden until it failed a key test about a year before launch. A significant effort was mounted to solve the problem and the decision was made to drop the original system and go to a new one which was completed within a year. This decisive action did demonstrate that it was possible to make major positive changes and move quickly, when it was necessary, and achieve long term flight operation efficiencies through commonality across multiple missions.

<u>Implementation</u>

7. <u>A strong systems engineering capability is needed for large, complex system development.</u>

Although this seems an obvious “motherhood” statement, this was not the case on EOSDIS despite the availability of systems engineering talent. Strong systems management, which is the combination of project management and systems engineering, is necessary for a successful development. The project manager should insist on it and the systems engineering leadership should support the project manager. The weaknesses in the requirements management and change control processes demonstrated that this was not the case.

[16] One downside to having the FOS developed by the flight segment developer is the difficulty of getting different developers to develop compatible systems operated in a common control center. This is a valid concern, and one that can be addressed by distributing appropriate specifications from the common control center developer.

These weaknesses were recognized at Hughes in late 1996 and 1997 by the incoming manager. He hired systems engineers experienced in large data systems engineering. At this point, along with several other changes, the government and contractor systems engineers started working together. One example was putting more emphasis on interfaces with the users.

8. <u>If the underlying processes are not in place, you don't have a chance.</u>

Processes like configuration management are basic necessities for running a project. A Hughes manager who joined mid-project observed that these processes "didn't exist at Hughes". They were weak on the government side as well. They had engineering change proposals that were 3 to 4 years old, and they had never been acted on. When change began to happen and positive changes were proposed, the system was not effective in its response; the proper tools are needed for estimating and understanding costs. A good trade between functions and costs is needed.

9. <u>Program, Project, and executive leadership must be aware of the environment. Had it been recognized, a technology paradigm shift that impacted the EOSDIS development would have changed the development path.</u>

During the procurement process, the World Wide Web was coming into its own. The concept of a centralized system using 80s technology was being challenged before the ink was dry on the contract. The project did adopt a technology evolution model, but scientists claimed it was too conservative.  In fact, even the model that the scientists proposed, which was considerably more aggressive than the project model, was too conservative. The technology price/performance curve was not being used in favor of the project.

Some new architectures were studied and proposed by the science community. It is not clear whether or not they would have been successful if attempted. However, at that critical point when technology forced things to diverge, it would have been useful to evaluate the direction of the project. It would have taken strong leadership, but at that point a change in direction would still have been possible.

<u>Recovery</u>

10. <u>Strong leadership, at all levels, is critical for development of a new, complex, highly-visible system.</u>

In the development of a complex, long-term project, the nature of the work changes as it moves from initial vision, through formulation, into implementation, and operations. It is crucial for the leader(s) to understand the system and its objectives, to direct a team through the complex procedures of development, to keep the team focused, to keep expectations under control, and to adjust to changes in the environment.

A lack of experienced leadership at all levels contributed to the early development difficulties of the EOSDIS. Although the vision for the early EOSDIS came out of Headquarters, there was a lack of program management experience in this area. This was critical, because the large program required direction and communication across diverse interfaces with many agencies. The GSFC also suffered from a lack of experienced project managers in large data systems, especially when compared to the expertise available in flight systems. The science community was not unified. Individuals in that community were not experts in the new technology or in the contracting area in which they were proposing change. Similar conditions also existed in the contractor's organization.

The first and primary ingredient for successful systems management is a strong government leader, or project manager. There were a number of times during development that were later recognized as points where a strong stand could have precluded further divergence in the project. It was clear at the SRR that there was a difference between how the contract was actually structured and how some believed it should be structured. This major disagreement could have been faced and acknowledged; the "head butting" could have been managed. A strong leader could have said, "OK, we need to change the contract dramatically," or "How can we be successful within the existing resources, using this contract, and with the understanding that we will all have to compromise?" This would not have been easy, but it was possible.

When things finally started to turn around, it was because the AA accepted the project's recommendation to go with SIPS, limited the products, and worked to protect the project. The later project managers also limited the requirements, established structure within the project, and involved the contractor as partners. The contractor brought people who understood systems engineering, worked with the mid-level mangers as partners, and went along with a restructuring plan that (although it removed work which went to the SIPS) allowed an end-point to be reached. These examples of strong leadership were essential for the project's turnaround and successful completion of the EOSDIS.

11. <u>Maintaining partnerships between the teams is necessary for a successful development.</u>

Partnership between teams is crucial to the success of a project, especially a large one. When issues arise, the attitude cannot be, "This is your problem, fix it." Rather, teams have to say, "This is our problem; how do we fix it together?" This does not negate the fact that contractual responsibilities must be met, but instead stresses that the success of one project team depends on the success of them all.

Unfortunately, there were times when the science community lost confidence in the project. There was not a "buy-in" from that community, and the government and contractor viewed each other as enemies. When the going was very tough, some contractors saw themselves as victims who were constantly harassed. It takes strong leadership to maintain partnerships under such adverse conditions. A strong leader is one who understands the system, has confidence in himself and his team, can be tough when necessary, but knows how to compromise.

These qualities emphasize creativity, along with a willingness to listen. For example, the contractor must succeed for the government project to succeed. This could mean restructuring a contract to provide a reasonable fee in the future, even if it is provisional. This provides an incentive which strengthens the team's commitment to project success.

One of the main factors in the turnaround was the strengthening of the government/contractor partnership. The leadership sought to improve this important relationship.

12. A large government program with high visibility draws political attention that can impact development.

This is often unavoidable for systems like the Shuttle or Space Station. However, it was not necessarily unavoidable for the EOSDIS. One of the original goals was to have a large, centralized system. This system would correct past abuses and attract a large defense contractor with the necessary "classified world" experience to do the job. In hindsight, there is compelling evidence that an incrementally developed, decentralized system would have been more effective in the long run.

Contractors, especially ones with a large workforce, are good lobbyists. This is another reason to maintain the previously discussed partnership. When problems are described by all teams as "our problem," Congress can be briefed by both the government and contractor working together.

13. Endless reviews do not help a struggling project.

Because of the amount of money, the large workforce, the number of people talking directly to the Administrator and Congress, and because of the difficulties encountered, it was natural for the project to attract a lot of attention. One consequence of this was the extraordinary number of external reviews requiring project response. This added even more to the project's already-crowded plate. These reviews were conducted by the Government Accounting Office, the Inspector General, NRC, and NASA, in addition to the normal project reviews. The project averaged one major review a month between 1993 and 1997, many of which resulted in conflicting recommendations. Thus, along with their regular work, the ESDIS Project was continuously preparing for a review and answering actions from the previous one.

The reviews themselves did not cause the turnaround. This was done by determined leadership within the project and at Headquarters. The leadership was instrumental in limiting the number of reviews so that the recovery process could begin.

14. Comparison of the ESDIS Project with the Jack Lee Study factors.

In 1992, a NASA study headed by former Director of the Marshal Space Flight Center (MSFC) Jack Lee, looked at factors that drive NASA program costs and technical risks. Eight factors were identified. These factors and their corresponding examples within the ESDIS Project are shown in the following table.

| **Major factors that drive NASA program costs and technical risks**. | **Assessment for ESDIS** |
|---|---|
| 1. Inadequate Phase B definition | The amount conducted was adequate for a typical project using an end-item spec. However, the issue of which development approach to use was not adequately covered and decided. |
| 2. Unrealistic dependence on unproven technology. | A more accurate description for EOSDIS might be: inadequate consideration of technology improvement and technology insertion in the program planning.<br><br>This was a problem. There was a mismatch between the EOSDIS build releases (2-3 years) and the technology lifecycle (6-9 months) and this mismatch was never resolved. Budgets were reduced based on new technology assumptions. These assumptions did not always pan out.<br><br>Also, it was unrealistic to assume that the application of object-oriented software development methodologies (which were emerging in the early to mid 90’s) could be applied against the same program baseline that was bid on a central mainframe-based approach. |
| 3. Annual funding instability. | This was a major problem. There were major cuts in the budget, most without a corresponding reduction in requirements. |
| 4. Complex organizational structure, including multiple unclear interfaces. | This was a problem. The science community was very diverse and not of one voice. The program structure and the project organizational location changed many times. The project was left without a strong Headquarters champion to decide between competing recommendations. |
| 5. Cost estimates that are often misused. | This was a problem. The initial funding looked to be adequate with sufficient contingency. But the money was distributed differently as disagreements |

| | |
|---|---|
| | concerning development philosophy cropped up between the project and the science community. An attempt to associate costs and functions was made (which should have been possible) in order to restructure. This couldn't be resolved, primarily because of the monolithic system design of the ECS, so there was never an agreement on how costs and functions were related. |
| 6. Scope additions due to requirements creep. | Some requirement expansion was expected, because of the evolutionary nature of the system. Technology changes would take care of the expansion. However, until the requirements were put in a box with Option A+, there was an unrealistic and impractical attempt to cover too many functions in one large system. |
| 7. Schedule slips. | This problem was due to factors like the underestimation of the amount of required code, skill shortages, rapidly changing technology, and inadequate processes. |
| 8. Acquisition strategy that does not promote cost containment. | The acquisition strategy was flawed and not suited for a large IT project. It provided for a rigid, end-item development instead of a more flexible, concentric development. |

**CONCLUSIONS**

The EOSDIS is now an operational system. It is handling large amounts of data for the Earth Science Enterprise and continues to add to its data base as new missions come on line. There are over two million users per year, far more than the 10,000 scientists who were originally expected to utilize the system. Ninety percent of these users are not scientists. There have been reports that  NASA data is more accessible today than it once was. This was one of the original goals of the EOSDIS. It is a major NASA system and will help scientists in their study of Earth for years to come.

There are still complaints that the system is "overly bureaucratic," i.e., not user friendly. It has neither all the functionality originally envisioned, nor all the user services originally planned. The reprocessing capability is lower than planned, which means reprocessing time is longer. These shortcomings are the results the EOSDIS rough development, caused by internal development mistakes and a changing external environment. As a consequence, the EOSDIS development had to change, and these changes impacted the final design. The "doable" EOSDIS could not deliver all that was originally promised. Had some of the changes been implemented earlier--such as focusing on a small doable core and expanding outward to include more innovative concepts like the SIPS--some difficulties could have been avoided. The externally-driven funding reductions also had a major impact.

The EOSDIS was conceived and started in one NASA era (1980s) and is being completed in an entirely different one. The fact that it is being completed at all, despite the internal and external problems, is a tribute to the tenacity of those who persevered to its completion. A number of lessons can be drawn from its development. If these lessons had been learned sooner, they could have contributed to a smoother development. Nevertheless, sharing these lessons can hopefully ease the development of future programs and projects.

## References


1. “Report of the EOS Data Panel, Volume IIa, Earth Observing System Data and Information System,” NASA Technical Memorandum 87777, 1986 (“Black Book”).

2. “Execution Phase Project Plan for Earth Observing System (EOS),” 170-01-01 Revision A, Goddard Space Flight Center, May 1995.

3. “EOSDIS Alternative Architecture,” Project Sequoia 2000 Team, Final Report, Contract#ECS-00012, Submitted to HAIS, September 6, 1994, Revised January 23, 1995.

4. “Special Review of the GSFC Earth Science Program/Project Management”, (Littles’ Review), presentation by the ESDIS Project, GSFC, June 18, 1998.

5. “Special Review of GSFC Earth Science Project Management,” Wayne Littles’ Team, June/July 1998.

6. GSFC Response to Special Review of GSFC Earth Science Project Management, July 30, 1998.

7. “Moving Forward: Lessons Learned,” excerpted from an ESDIS Project presentation to the National Research Committee, summer 2001.

8. “A Review of the U.S. Global Change Research Program and NASA’s Mission to Planet Earth Earth Observing System”, National Research Council, Washington D.C., 1995.